mpEAd: Multi-Population EA Diagrams

Sebastian Lenartowicz[1] **and Mark Wineberg**[2]
University of Guelph
50 Stone Road, Guelph ON N1G 2W1
[1]*slenarto@mail.uoguelph.ca*
[2]*mwineber@uoguelph.ca*

arXiv:1607.05213v1 [cs.NE] 18 Jul 2016

Abstract— *Multi-population evolutionary algorithms are, by nature, highly complex and difficult to describe. Even two populations working in concert (or opposition) present a myriad of potential configurations that are often difficult to relate using text alone. Little effort has been made, however, to depict these kinds of systems, relying solely on the simple structural connections (related using ad hoc diagrams) between populations and often leaving out crucial details. In this paper, we propose a notation and accompanying formalism for consistently and powerfully depicting these structures and the relationships within them in an intuitive and consistent way. Using our notation, we examine simple co-evolutionary systems and discover new configurations by the simple process of "drawing on a whiteboard". Finally, we demonstrate that even complex, highly-interconnected systems with large numbers of populations can be understood with ease using the advanced features of our formalism.*

1. Introduction

From the beginning, it has been obvious that evolutionary algorithms (EAs) could make use of multiple populations in order to facilitate more complex searches and increase the power of the search itself. Though it is doubtless that others exist, there are four main types of multi-population systems that have been investigated. In the island model[1], solutions move between different discrete populations that use the same objective function. The predator/prey model[2] uses multiple populations to perform fitness evaluation – an individual in one population is compared against one or more individuals in another population, where its fitness score increases as those of the others decrease (and vice versa). Co-operative co-evolution[3] is yet another different system, in which members of each different population comprise different elements of the complete solution (and members must be drawn from each population to form and evaluate a full solution). Finally, hierarchical systems, though sparsely investigated (and with conflicting definitions[4][5]), use a variety of different multi-population structures that utilize levels in order to achieve their aims.

Investigation into multi-population EA systems has waxed and waned over the decades, likely because such systems tend to produce complexities that simple single-population systems do not incur. Populations may exchange both genetic and evaluative information[1][6], and, in more esoteric systems, other types of information as well[7]. There are often complex spatial structures connecting the populations formed by this web of relationships[1], and the structure may be distinctly different depending upon the nature of the information[7]. Furthermore, there are recursive effects between multiple populations – as seen in many co-evolutionary systems[3][6][8] – that lead to problems such as Red Queen[9], the loss-of-gradient effect[8], and decoupling[8]. This complexity in the structure of information flow throughout the system is in addition to the actual movement of individuals between populations, as seen in island-model migration[1]. All of this is further exacerbated when considering hierarchical multi-population systems. Confusion about the very idea of what constitutes a hierarchical system can easily be seen in even a cursory review of the literature[5][7][10]; all of these systems are called hierarchical and incorporate elements of hierarchy, but all are different, with very few elements in common!

To combat the confusion arising from all this complexity, we have developed a graphical formalism that encapsulates the different relationships that can exist between multiple populations in an EA system. The multi-population EA diagram (mpEAd, pronounced "emm-pede") employs a concise visual grammar to depict multiple populations and the information flow between them, in a similar way to how the Unified Modelling Language (UML)[11] captures the categorization of and relationships between the component parts of object-oriented software systems.

At this juncture, it is important to note that the mpEAd system is not intended to depict the internal mechanisms and dynamics of single populations; in other words, it does not describe the different selection methods, reproductive operators, etc., of the evolutionary algorithm used by a population. Instead, mpEAd treats each population as a black box, accepting inputs and producing outputs for consumption by other populations or itself. This is approximately analogous to the static "structural diagrams" of UML, which depict the relationship between the system components (such as classes and subsystems) rather than their time-dependent activities, which are modelled in its "behavioural" diagrams. It should also be noted, however, that, while inspired by UML, mpEAd does not incorporate any of the notational conventions found in it, instead using a visual language that is more suited to modelling EA systems. Finally, it is important to stress that mpEAd is more than the topological structure of the populations, as frequently seen when discussing migration. The relationships depicted in mpEAd are much broader in scope, and model all types of information exchange between populations, with migration being only a single subset.

This paper is divided into two parts: the first presenting the basic elements of mpEAd necessary to structure any system and the second extending this notation to model more complex systems. Examples are provided throughout to demonstrate the notation and power of the mpEAd system.

2. Essential Structure of mpEAd

All attempts to depict multi-population EAs base themselves on the graph formalism, with populations represented as nodes (with edges taking variable meanings). However, as with class diagrams in UML, where nodes are classes and edges the relationships between them (messages, roles, etc.), the complexity of an EA system is not fully captured by a simple graph and additional diagrammatic formats are required in order to capture all aspects of the EA's functioning.

With this in mind, a few design principles emerged organically while constructing the mpEAd formalism: intuitiveness, consistency, distinctiveness, and simplicity. While the roles of intuitiveness and simplicity are simple and intuitive (and yet sometimes difficult to achieve), a brief explanation of the other principles is warranted. While consistency may seem similarly straightforward, the notion

of establishing a visual grammar is often lost when attempting to communicate information. Ensuring that similar things appear similar helps bring to mind meaning and allows ease of learning, extensibility, and creativity. In contrast to this, differences should appear different. This promotes readability and ease of interpretation. We also held to two supplementary principles: that the diagram should render naturally in black and white (for publication), and that it should be simple to draw by hand on a whiteboard or piece of paper.

2.1 Basic Elements of mpEAd

The graph at the heart of mpEAd incorporates two types of nodes: population nodes and computation nodes, as well as a number of different edge types. All of these are described in detail below.

2.1.1 Population Nodes

A population node corresponds to a single optimization algorithm (usually an EA) and a set of solutions (the population). It is denoted using a simple hollow rectangle with solid borders and the name of the node written inside. The population node, when drawn, always includes a set of multiple parallel lines, which serve as a visual reminder of the many members of the population inside.

2.1.2 Information as Edges

Edges in the mpEAd graph are used to model the information flow between nodes. As information flow is directional, mpEAd becomes a directed graph, and, per the convention, uses arrows to indicate direction. The types of information available in a multi-population EAs are numerous and varied but can, however, be categorized into two distinct groups: genetic and non-genetic. Genetic information consists of any information used to construct or embody a solution, being referred to as genotypic and phenotypic, respectively. These are often interchangeable, as the genotypic is often immediately evaluated for fitness as if it were a phenotype. Furthermore, even when phenotypes are used, they are produced and consumed during a single decoding/evaluation step, and discarded upon completion. This often stops being true in multi-population systems, where phenotypes may be passed around for various purposes and used by multiple evaluations. The flow of genetic information is represented using a solid edge with a closed arrowhead that is either hollow (genotypic information) or filled (phenotypic information). Examples of these edges are given in Figure 1.

Evaluative information, such as objective and fitness values, is nearly universal in optimization; this information is represented in mpEAd using a dashed edge with an open arrowhead, with the dashed edge serving to indicate that evaluative information is non-genetic in nature. While other types of non-genetic information, such as statistical or control information, may exist in multi-population optimizing systems, a full discussion of these is outside the scope of this paper.

If the same unit of information has multiple simultaneous recipients, the edge is drawn as diverging from a single source. Conversely, if two pieces of information are required, the lines to the node are kept separate and distinct and are not combined in a similar way.

2.1.3 Computation Nodes

The second type of node, the computation node, is less obvious and is one of the elements that makes the mpEAd formalism more than simply a topological model of the connections between populations. The role of the computation node is to take in one or more streams of information, perform processing on them, and to provide the result to another node or nodes. Computation nodes perform a variety of information processing operations, including but not limited to decoding genotypic information into phenotypic information, evaluating fitness, and combining information from different sources. The computation node is depicted using a large hollow circle, often labelled with a name, such as the name of the fitness function used for evaluation.

In general, the border of the circle matches the line type of the output edge(s). If the output is of mixed type, the circle's border alternates evenly between solid and dashed. It should be noted that the fundamental difference between a computation node and a population node is that, while both types can perform information processing, the computation node is stateless and does not store information, only taking input and producing output based upon it.

Depiction	Type of Information
	Phenotypic
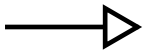	Genotypic
	Evaluative

Fig. 1: Arrowheads and edges used in mpEAd.

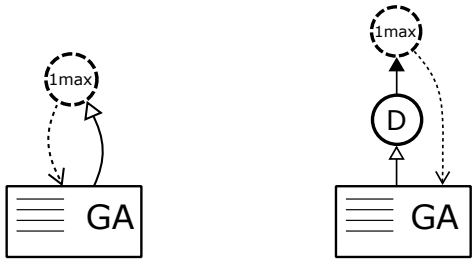

(a) A basic single-population EA solving the 1max problem.

(b) A variant that decodes genotypes into phenotypes for evaluation.

Fig. 2: Examples of a basic EA.

2.1.4 Putting the Basic Elements Together

The simplest multi-population system is one with only a single population; i.e. the simple EA. Two simple examples of this kind of system, solving the universally-known 1max problem, can be seen in Figure 2 to provide context for understanding the basic elements of mpEAd. In Figure 2a, the genetic information is evaluated directly, with the resulting fitness value being passed back to be stored within the population. In Figure 2b, the same evaluation takes place, but must first be decoded into a phenotype before evaluation can take place.

2.2 mpEAd in Action

While efforts have been made in the past to model the interactions between populations, they are often simplistic and rely on *ad hoc* notations, similar to what is seen in Figure 3. The power of mpEAd becomes apparent in comparison to this, as it permits much more accurate and detailed modelling of how the populations interact. All of the diagrams in Figure 4 are different co-evolutionary systems that

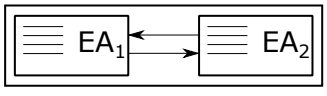

Fig. 3: A typical naïve way to model the systems in Figure 4.

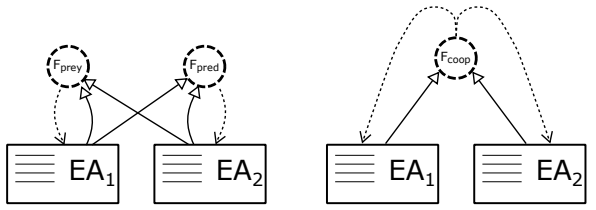

(a) A basic predator/prey co-evolutionary system.

(b) Standard co-operative co-evolution.

(c) A hybrid of co-operative and predator/prey co-evolution with modification.

(d) A hybrid of co-operative and predator/prey co-evolution with cross-modification.

Fig. 4: Examples of various co-evolutionary systems. The two complementary systems in 4c and 4d are both, to the best of our knowledge, novel.

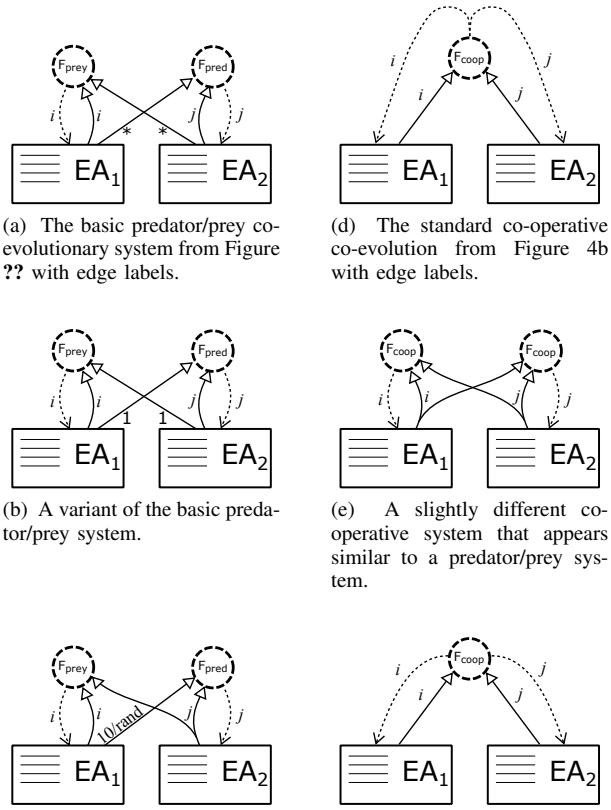

(a) The basic predator/prey co-evolutionary system from Figure **??** with edge labels.

(d) The standard co-operative co-evolution from Figure 4b with edge labels.

(b) A variant of the basic predator/prey system.

(e) A slightly different co-operative system that appears similar to a predator/prey system.

(c) A more complex variant of the predator/prey system.

(f) The same system as depicted in 5e that is more clearly co-operative (though still subtly different from 5d).

Fig. 5: Examples of edge labels.

would be equivalent to the one in Figure 3. Many disparate types of multi-population systems (in this case, a variety of co-operative and competitive co-evolutionary systems) can thus be represented in a way such that their similarities, as well as their differences, become apparent.

Figures 4a and 4b depict a pair of standard co-evolutionary systems that are familiar to most EC researchers. In Figure 4a, the diagram shows both EA populations sending members to the predator and prey evaluation functions, which are used to compute the two different fitnesses. Figure 4b, in comparison, depicts a co-operative system – where the individuals from the two populations are combined to produce a single fitness value that is applied to both.

With the co-evolutionary systems in Figures 4c and 4d, interesting possibilities begin to appear. These systems are unknown in the literature, but by using mpEAd can easily be conceived of, modelled, and constructed. On examination, they appear to be a hybrid between the co-operative and predator/prey systems seen in Figures 4a and 4b. For both of these systems, a single fitness value is produced based on input from the two different populations; however, for one of the populations, the fitness value is modified, either by the individual itself or the individual from the other population. Many different co-evolutionary systems can be easily constructed in an analogous manner, demonstrating the power of the mpEAd formalism for both modelling and discovery.

3. Advanced Features of mpEAd

While mpEAd has many additional, powerful features that make modelling even very complex systems trivial, there are too many to exhaustively discuss here. Instead, we concentrate on the ones necessary to provide the most understanding for the most systems. To this end, we describe four additional features: edge labels, inset arrowheads, macro boxes, and ellipsis notation.

3.1 Edge Labels

The co-evolutionary examples given in Figure 4 hide a great deal of important detail regarding the structure of the information being passed around the system. In particular, when considering evaluation, they do not provide information about how many individuals are required, where they come from, how they are to be selected, where they are to be stored, etc. There are many approaches for matching individuals between co-evolutionary populations, ranging from simple pairing, to pairing with an elite, to exhaustive combination pairing. Because these different constructs would result in an mpEAd that otherwise looks the same, edge labels can be used to disambiguate the selection and matching of individuals between nodes.

The simplest kind of edge label is a letter variable, which is used to indicate sequential iteration through the individuals in a given population, both for selection and storage of incoming values. These variables can be thought of as indices to individuals within the population. Similarly, numbers (either single or in a range) are used to indicate when multiple individuals are drawn from a population in order to perform a computation. The algorithm by which these individuals may be chosen can be specified by a forward slash followed by an algorithm name or symbol (which should be described in accompanying literature) following the number or range. A more thorough discussion of this notation is outside the scope of this paper and will be explored in future work. An asterisk (*) is a special case of numeric value, in which the entire population is used and, in this case, no algorithm specifier may be provided. The asterisk was chosen rather than the more conventional n used in computer science because n could be mistaken for an index variable, and the asterisk as a symbol is commonly used to denote "everything" or "all" (e.g. as used in string matching and the Unix command line).

It should be noted that these categories of edge labels were developed while considering currently existing co-evolutionary and multi-population systems, and are likely to be far from exhaustive. It is almost unquestionable that extensions to the edge label notation will occur in future work as more use cases are considered.

3.1.1 Examples

The utility of edge labels becomes apparent when considering the mpEAds in Figure 5, all of which are some variation on two-population co-evolution (as seen in Figures 4a and 4b).

Figures 5a to 5c describe variations on Figure 4a. The first example, Figure 5a, is a common implementation of the predator/prey model, in which each predator is tested against all prey, and each prey is tested against all predators for their respective fitness values. This is computationally expensive (being $O(n^2)$), and subsequently, variants using less than the full population are common.

Figure 5b represents such a system, in which each individual from one population is paired with some other individual from the other population for evaluation. The method by which the other individual is selected is left unspecified. This individual could be randomly chosen, be the most fit, or selected by some other algorithm.

In Figure 5c, which represents an asymmetrical approach to predator/prey (used only for demonstration purposes – so far as we are aware, no such system exists in the literature). In this example, each predator is compared against ten prey, using the specified selection mechanism (drawing them at random). The prey, meanwhile, is evaluated in a fixed 1:1 pairing with a given predator.

Figure 5d depicts the simplest co-operative system that can be inferred from Figure 4b, using a similar 1:1 pairing of individuals from each of the populations to that seen in the predator/prey exaple in Figure 5c. This structure, though outwardly very different, is actually quite similar to that seen in Figure 5e, which, despite its superficial resemblance to the various predator/prey systems, is actually co-operative due to the 1:1 mapping between i and j for evaluation and the use of the same objective function for the evaluations.

There is, however, a very subtle difference between Figures 5d and 5e: in 5d, the same information is sent to both populations (being stored at locations i and j), whereas, in 5e, the information sent to both populations may not necessarily be the same. Finally, Figure 5f actually depicts the same functional system as 5e, as the different information streams are depicted separately rather than coming from the same source.

3.1.2 Examples of Greater Complexity

Figure 6a depicts a variation on predator/prey in which the raw genotypes are first decoded into phenotypes before being used for standard predator/prey evaluation. It should be noticed that the borders of the decoding functions are solid, as the output of the computation nodes is phenotypic (and therefore genetic), which is represented using solid lines. The predator requires information from both the prey population and the predator population to be decoded before it can be evaluated, while the prey population only requires the predator to be decoded, while the genotypes of the prey themselves are acted upon directly.

In Figure 6b, a third population is introduced to model a more general version of predator/prey based on the work of de Boer, Folkert, and Hogeweg[12]. The third population is composed of "scavengers" whose fitness is dependent on the fitnesses of both predator and prey, but which do not affect the fitness of either. The scavengers, much like their biological counterparts, exist only to "pick up the scraps" after the predator and prey have finished

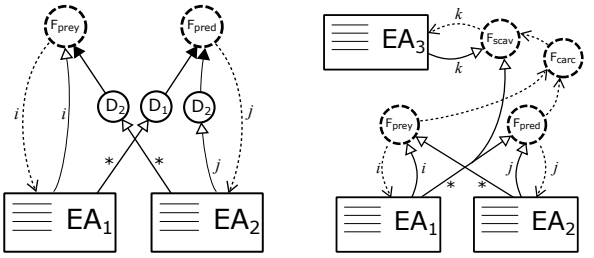

(a) A variation on predator/prey in which genotypes are decoded into phenotypes before being considered together.

(b) Predator/prey extended to include scavengers[12].

Fig. 6: Examples of more-advanced systems using edge labels.

evaluating each other. In this system, the scavengers require the prey genotype to establish "edibility"; after all, if the scavenger cannot digest the prey, it will go hungry!

In this example, in addition to the three objective functions (one per population), there is a fourth function, F_{carc}, where "carc" is an abbreviated form of "carcass". This function takes the evaluated fitnesses of each predator/prey pair and produces a "fitness value" for the consumption of the scavenger's objective function. This value, in effect, represents whether a given prey was actually killed (and can therefore be consumed), as well as how much of the prey is left over and can be used by the scavenger (hence "carcass", a term not used in the original paper[12]). The genetic information of the prey is also forwarded to the scavenger, in order to establish edibility as described above. It should be noticed that the mpEAd clearly represents all aspects of this process, as well as indicating the asymmetrical role of the scavenger population.

3.2 Inset Arrowheads

Migration between populations is a common feature of multi-population systems and it would be remiss to exclude it from mpEAd. Yet, migration presents an apparent quandary: at first glance, it appears to break with the "information flow" model, as an actual individual is being transported rather than formless information. On deeper reflection, it becomes obvious that migration can be modelled as a transfer of genetic information followed by a state change in the receiving population. There is, however, a distinction to be made between the arrival of genetic information to be added to the population versus the arrival of genetic information to be incorporated into existing individuals in that population. The first, obviously, is migration. The second is lesser-known, although still existing in the literature, where it is known as hierarchical composition[13]. In every other instance in mpEAd, arrowheads that touch a population node affect individuals within that population. Consequently, a genotypic arrowhead that touches the node border maps more closely to hierarchical composition (which changes the individual) than migration (which adds an individual to the population). Migration, then, is modelled using inset arrows located along the edge in question (as seen in Figure 7).

3.3 Macro Boxes and Ellipsis Notation

A macro box is, much like in a programming language, a simple shorthand for repeating elements. This is typically used in mpEAd to depict the same evaluation being used independently for multiple populations (this being distinct from a single evaluation shared by multiple populations, which implies co-evolution).

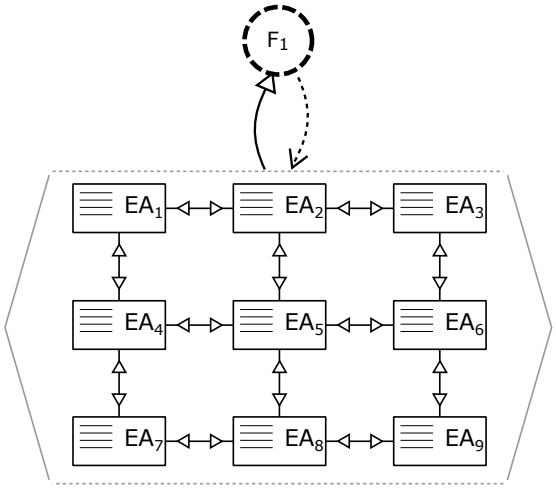

Fig. 7: A small island-model system using a 3x3 grid of populations.

An element that originates arrows ending at the border of the macro box is duplicated and connected to every element within the box, while elements whose arrows do not end at the macro box border are not duplicated. Macro boxes are are depicted using a elongated gray hexagon in which the long face is dashed and the remaining faces are solid. A simple example of macro box usage can be seen in Figure 7, in which a traditional island-model EA is shown with migration occurring across a 3x3 grid of populations. While macro boxes have many additional uses and properties, a full discussion of these is outside the scope of this paper.

The ellipsis (...) is used analogously to abbreviating text. It indicates the presence of many more identical components of a system, the actual drawing of which would be difficult or unwieldy. The ellipsis is always accompanied by a horizontal bar with vertical ends and an integer (value greater than 2) that indicates the total number of elements, including those actually drawn. An example of this usage is given in Figure 8, in which a 32x32 migratory grid is depicted in a compact form.

The full power of the combination of these two concepts can be found in Figure 9, which depicts an extended version of the hierarchical GA created by Sefrioui and Périaux[4]. This system is

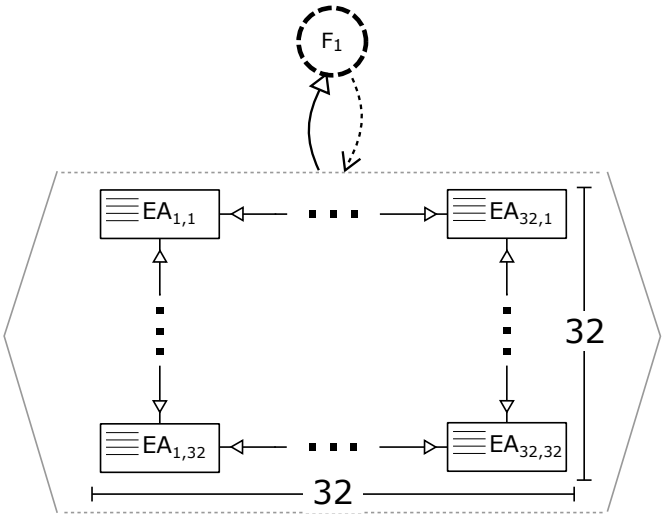

Fig. 8: A much larger island-model system, using a 32x32 grid of populations.

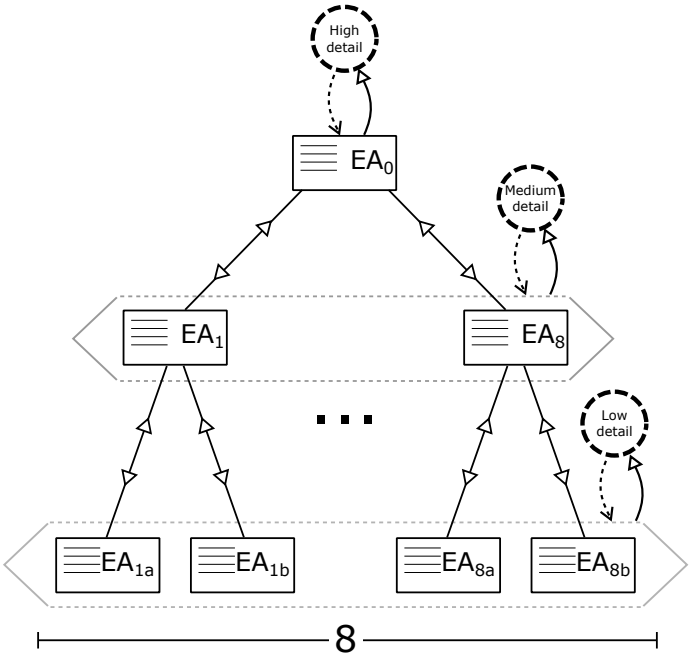

Fig. 9: An extended version of the hierarchical GA created by Sefrioui and Périaux[4].

used to find solutions for an problem for which the objective function is computationally expensive to evaluate, but can be approximated using functions that require less computational effort at the expense of precision. Sefrioui and Périaux first use a number of coarsely-evaluated populations to search for promising solutions, which are then passed upwards in order to be more finely evaluated. They also allow solutions to be passed down, in order to assist the lower-level populations and keep them "current". While the original version takes the form of a binary tree of populations, this variant adds additional subtrees below the "high detail" node, resulting in a system of 25 nodes rather than the original 7. Such a system would be able to more effectively cover the entire search space than the original model, potentially accelerating the process of locating a solution to a complex problem.

4. Conclusion

The mpEAd formalism is a graphical notation designed to permit the depiction of large, complex multi-population EA systems. Designed with the goals of being as intuitive, consistent, distinctive, and simple as possible, mpEAd is a powerful modelling tool for systems often considered too complex to describe clearly. Even with a system as small as two populations, we have seen that, through the use of mpEAd, it is possible to envision not only existing systems, but to diagram and reason about systems not in the literature in an easy and clear way. In the future, even more elements may be added to mpEAd as EC itself grows and matures, necessitating unenvisioned interactions and relationships. Ultimately, it could be possible to develop an "mpEAd IDE", where complex systems are created visually before being rendered down to source code automatically. While these pursuits remain in the future, it is clear that mpEAd has practical applications in the here and now, depicting the exceedingly complex in a simple manner.

References

[1] Darrell Whitley, Soraya Rana, and Robert B Heckendorn. "The island model genetic algorithm: On separability, population size and convergence". In: *Journal of Computing and Information Technology* 7 (1999), pp. 33–48.

[2] W Daniel Hillis. "Co-evolving parasites improve simulated evolution as an optimization procedure". In: *Physica D: Nonlinear Phenomena* 42.1 (1990), pp. 228–234.

[3] Mitchell A Potter and Kenneth A De Jong. "A cooperative coevolutionary approach to function optimization". In: *Parallel problem solving from nature-PPSN III*. Springer, 1994, pp. 249–257.

[4] Mourad Sefrioui and Jacques Périaux. "A hierarchical genetic algorithm using multiple models for optimization". In: *Parallel Problem Solving from Nature PPSN VI*. Springer Berlin Heidelberg. 2000, pp. 879–888.

[5] Mehmet Gulsen and Alice E Smith. "A hierarchical genetic algorithm for system identification and curve fitting with a supercomputer implementation". In: *Evolutionary Algorithms*. Springer. 1999, pp. 111–137.

[6] Kenneth A De Jong. *Evolutionary computation: a unified approach*. MIT Press, 2006.

[7] QS Li et al. "A multilevel genetic algorithm for the optimum design of structural control systems". In: *International Journal for Numerical Methods in Engineering* 55.7 (2002), pp. 817–834.

[8] R Paul Wiegand and Jayshree Sarma. "Spatial embedding and loss of gradient in cooperative coevolutionary algorithms". In: *Parallel Problem Solving from Nature-PPSN VIII*. Springer. 2004, pp. 912–921.

[9] Ludo Pagie and Paulien Hogeweg. "Information integration and red queen dynamics in coevolutionary optimization". In: *Evolutionary Computation, 2000. Proceedings of the 2000 Congress on*. Vol. 2. IEEE. 2000, pp. 1260–1267.

[10] Kim-Fung Man, Kit Sang Tang, and Sam Kwong. *Genetic algorithms: Concepts and designs*. Springer Science & Business Media, 2012, pp. 65–74.

[11] James Rumbaugh, Ivar Jacobson, and Grady Booch. *Unified Modeling Language Reference Manual, The*. Pearson Higher Education, 2004.

[12] Folkert K de Boer and Paulien Hogeweg. "Co-evolution and ecosystem based problem solving". In: *Ecological Informatics* 9 (2012), pp. 47–58.

[13] Richard A Watson. "Compositional evolution". PhD thesis. Brandeis University, 2002.